CoSine IP Service Delivery Platform Application Architecture

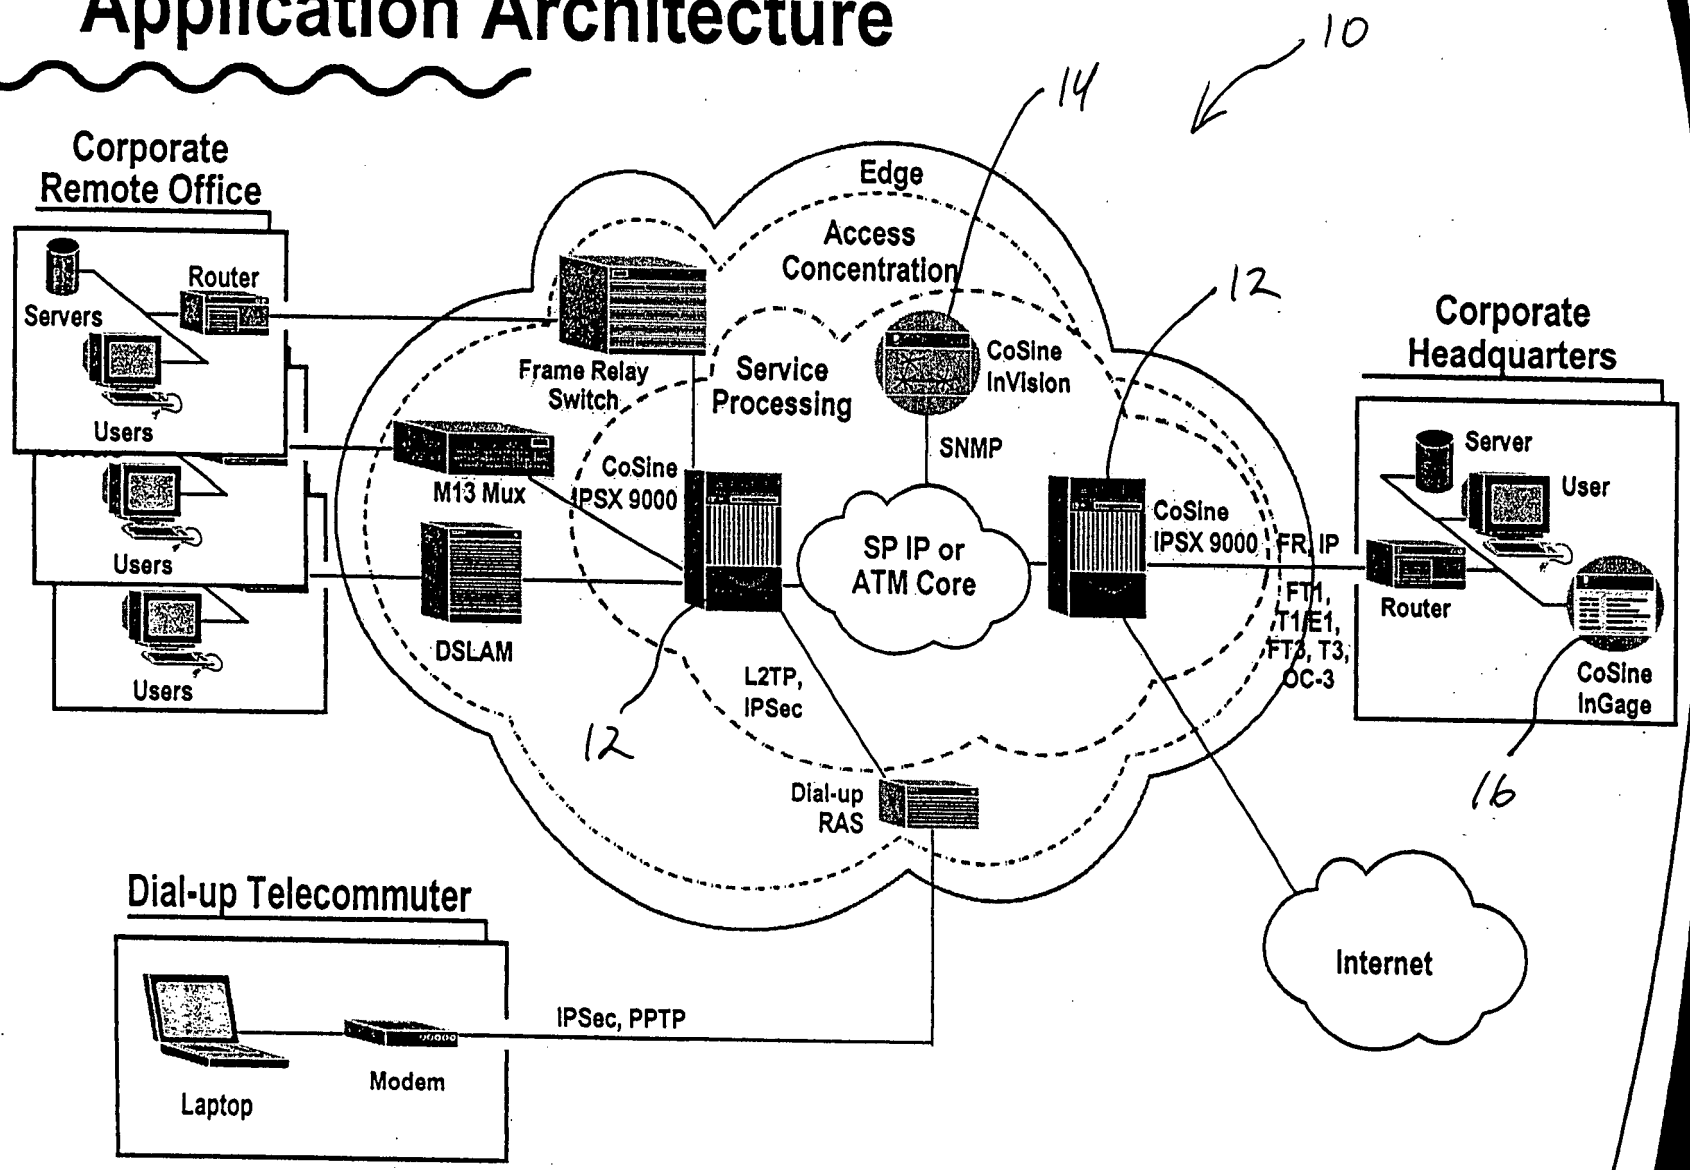

Diagram 5 — Managed Firewall Service with CoSine's Network-based Solution

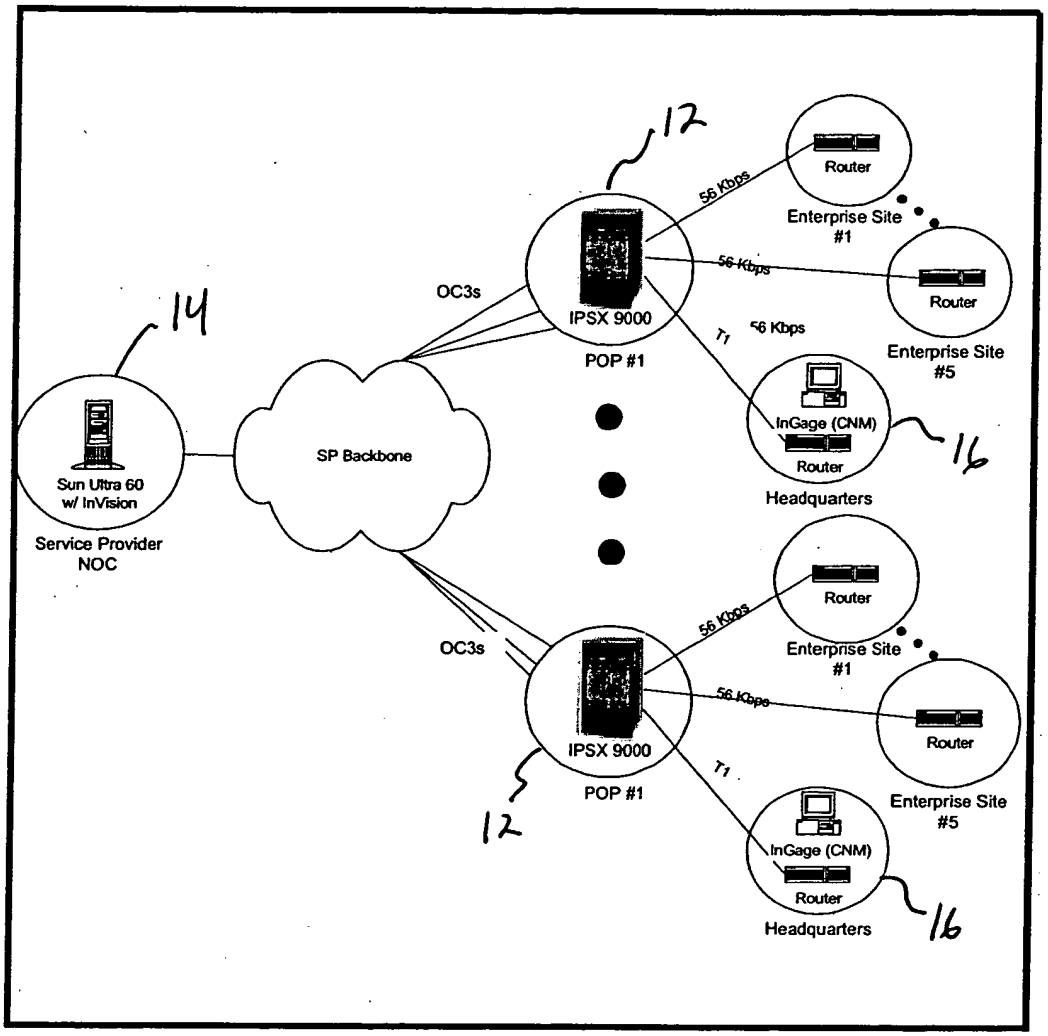

POP Infrastructure

The POP access infrastructure in the network-based managed firewall service model is based on the CoSine Communications IPSX 9000 Service Processing Switch. The base configuration for the switch includes:

- 26-slot chassis
- Redundant power supply
- IPNOS Base Software
- Ring Bridge & Ring Bridge Pass-Thru (to complete midplane)
- Control Blade (for communications with InVision Services Management System)
- Dual-port Channelized DS3 Access Blade
- Dual-port Unchannelized DS3 Access Blades
- Processor Blade
- OC-3c POS Trunk Blade

The following tables analyze the cost structure of all of the above models and projects these costs out over 5 years:

IPNOS Overview

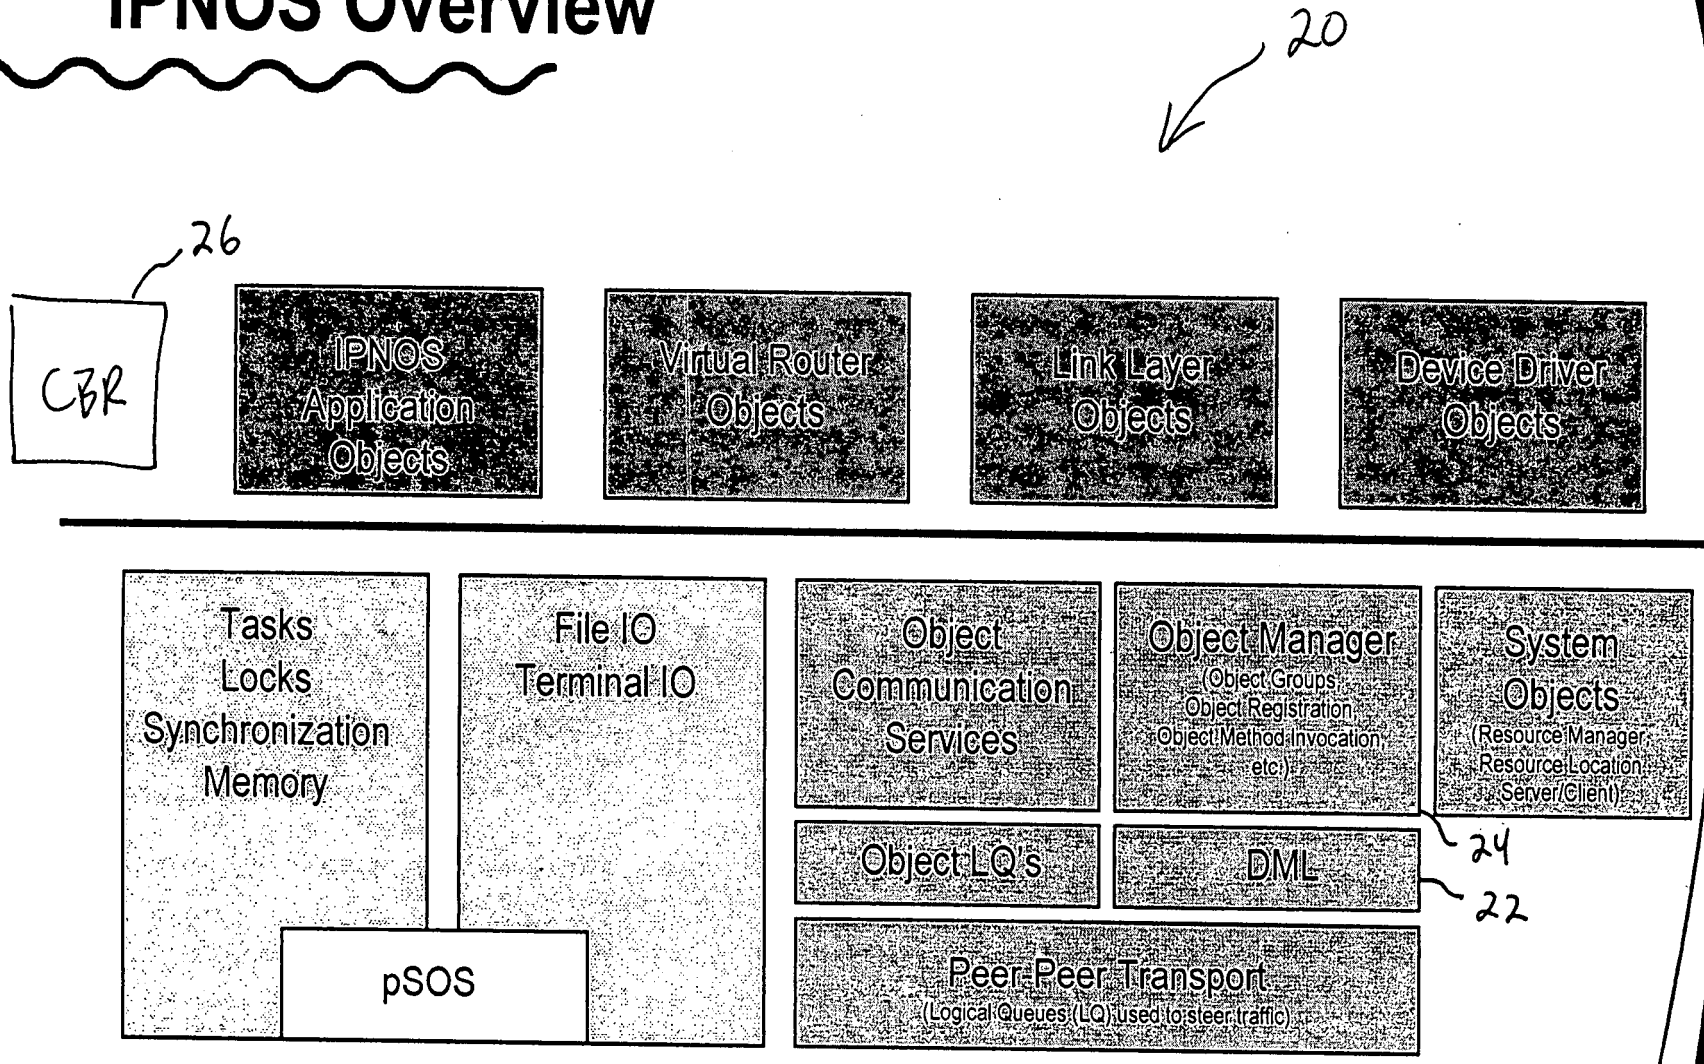

Fig. 3

CoSine COMMUNICATIONS

2. OM Design

2.1. Overview

The Object Manager consists of 3 layers as shown on Figure 1. The upper layer titled *OM Configuration Database* (OMCD) is concerned with managing the VPN and VR configuration. This is the agent that deals with CM directly. Middle layer titled is *OM Object Routing and Interface Global* concerned with managing global (across the IPSX system) object groups and object configurations. The lower layer titled *OM Object Routing and Interface* (OMORI) is concerned with managing local objects and groups as well as routing control information between address spaces based on the location of objects, and interfacing with the object via method invocation.

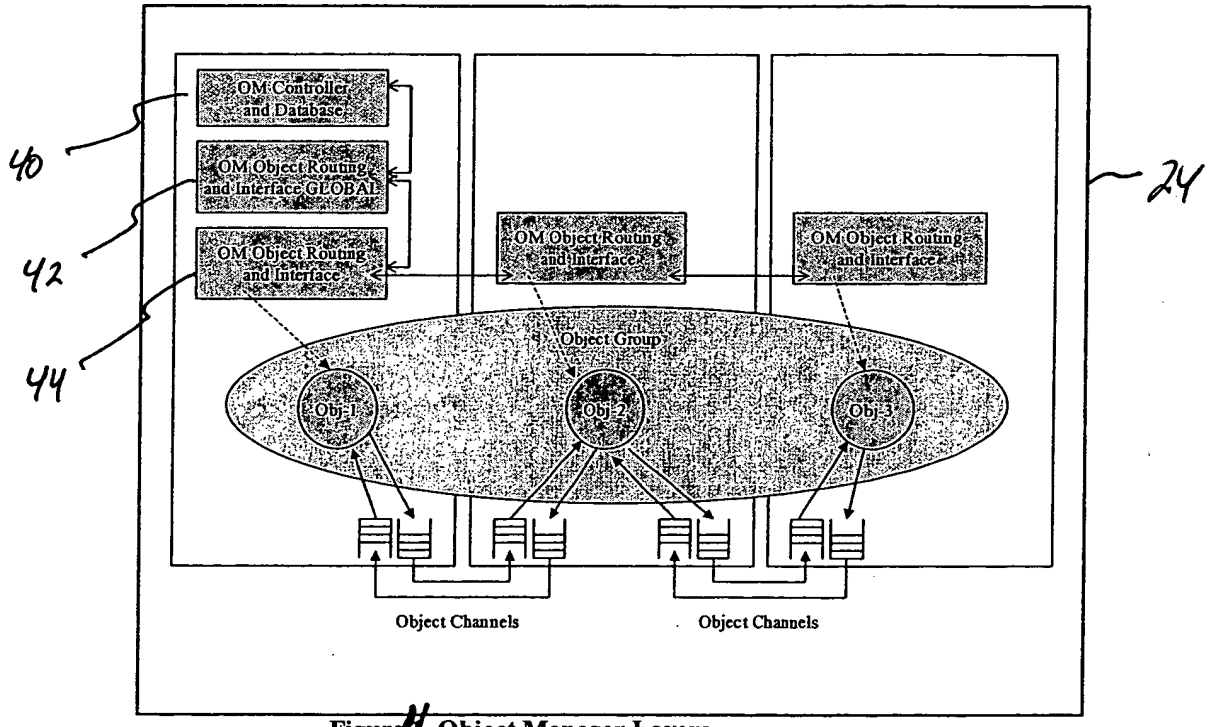

Figure 4. Object Manager Layers

2.1.1. Database distribution

IPSX object database consists of two types of database: Global, managed on Master Control Blade by OMORIG and distributed local databases, managed by OMORI agents on every PE present in the system. Global database is a superset of the extracts from local databases.

2.2. Object

Objects represent a basic unit of management for purposes of fault tolerance, computational load balancing etc. One or more adjacent protocol modules can be placed into a single object. It is also possible that a module is split across two objects.

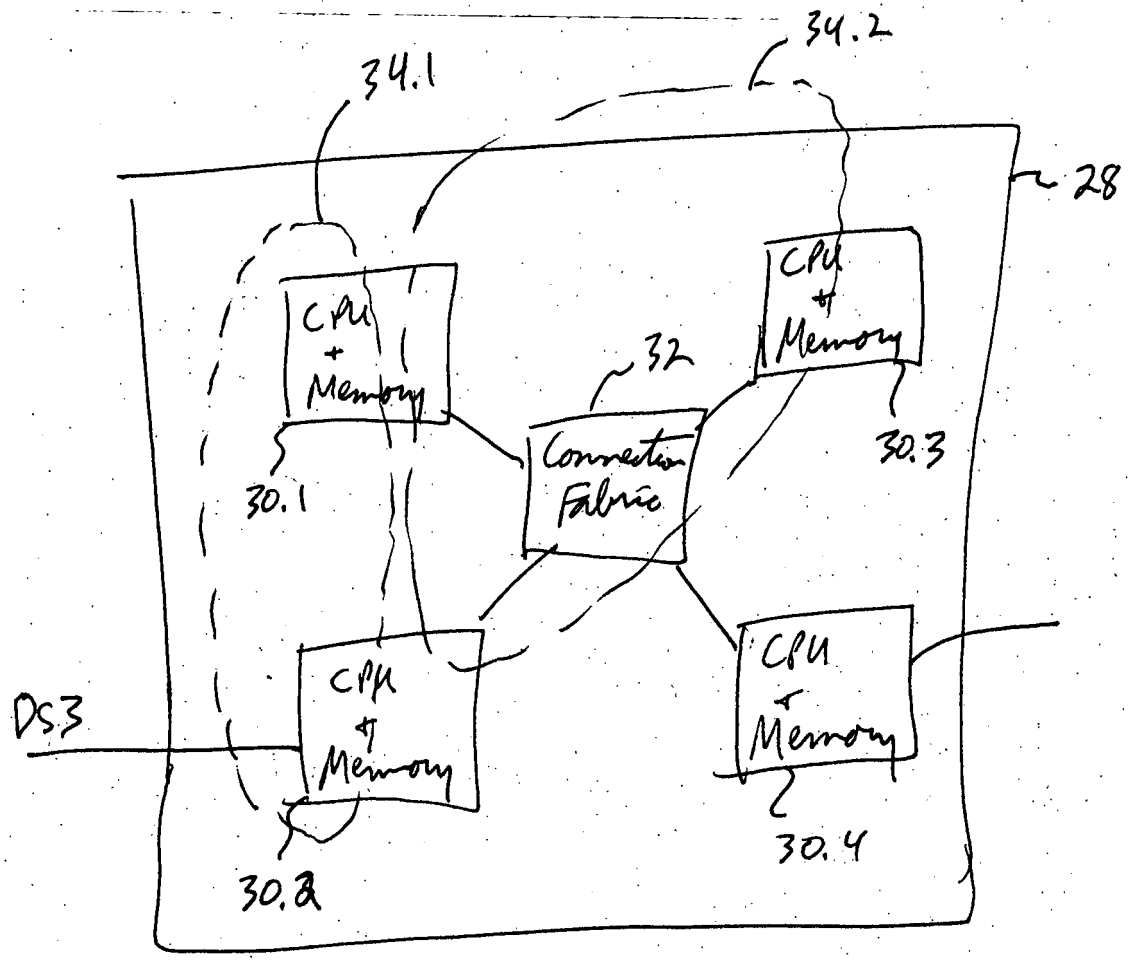
34.1
34.2
28
CPU + Memory
32
Connection Fabric
30.1
30.3
DS3
30.2
30.4

Fig. 5

Innovative Method for Segmenting and Layering Services

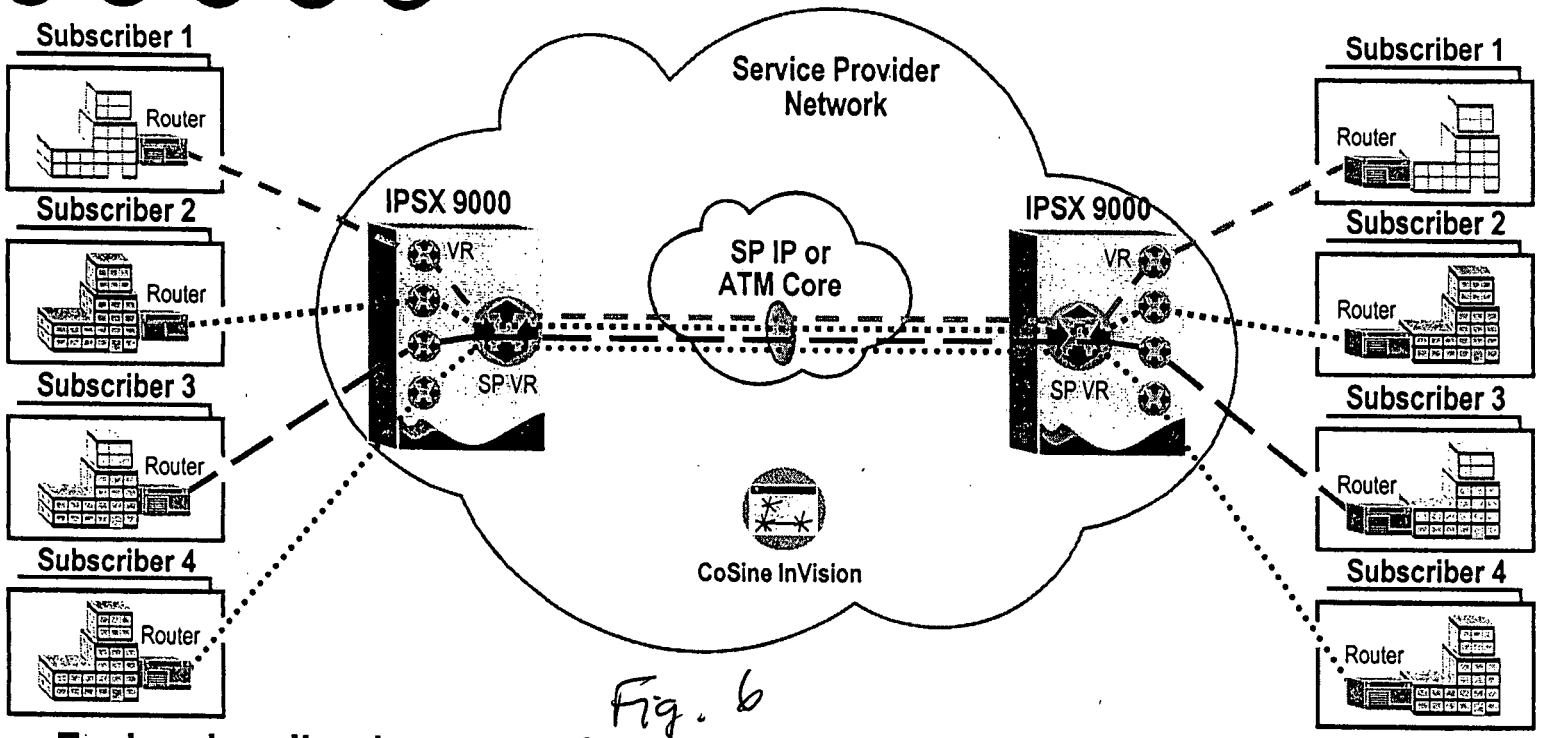

Fig. 6

- Each subscriber has a set of partitioned Virtual Routers (VRs)
- Each VR is the equivalent of an independent hardware router
- VR as an object group enables customized services per subscriber
- InVision allows ease of service provisioning and maintenance of services across all IPSX units in a SP network
- IP Network Operating System's (IPNOS) open Application Program Interface (API) enables new services to be continually added to the platform

CoSine COMMUNICATIONS

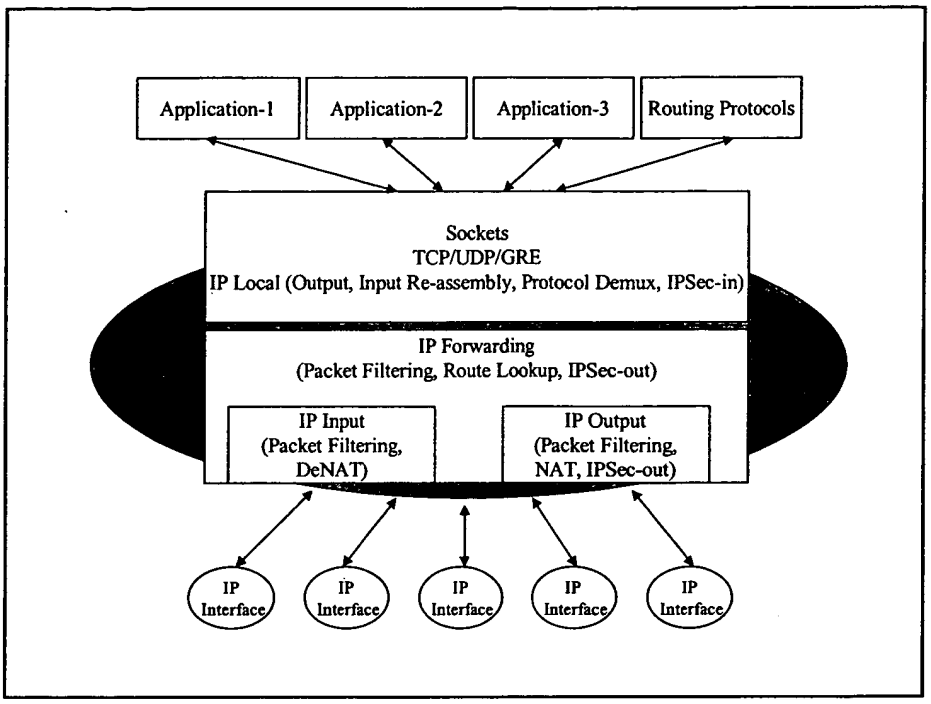

Figure 7: Standard Protocol Stack Profile

We now look at what would happen in IPNOS 1.x for typical application flows. The inter-module transfers are shown in Figure 8. The processing sequence is

1) Packet is received from the subscriber line interface
2) Processed by IP Fwd module and passed to IP Local module
3) Demux'ed by IP Local and passed on to the Application
4) Application process data and sends transformed data
5) IP Local sends packet to subscriber IP Fwd module
6) Subscriber IP Fwd forwards it to ISP IP Fwd module
7) ISP IP Fwd module sends it to ISP line interface

Note that steps 2, 3, 4, 5, 6 require an inter-PE transfer. Each of these steps may represent a backplane crossing.

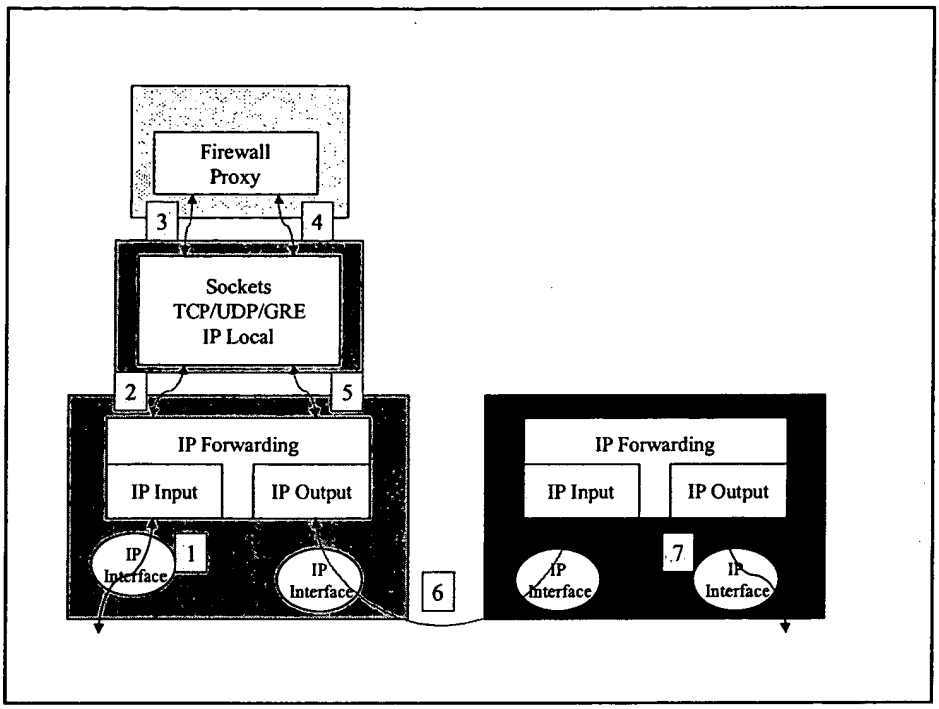

Figure 8: Firewall flow processing

2.3 Differentiating between L3 and L4 flows

Figure 3 shows the contents that can be used to determine a L4 packet flow when the packet is unfragmented (determined by the condition *(MF == 0 && Offset == 0)*). It also shows the contents that can be used to determine a L3 packet flow when the packet is fragmented.

When a packet gets fragmented, it should be noted that the ID field is preserved across all fragments of the packet. However, the receiver may receive the fragments in arbitrary order. Thus the receiver may only use <Source, Destination, Protocol> to determine the packet flow. The receiver can not create per packet flow state which uses the ID field to decide on the L4 flow for all fragments on the false assumption that the first fragment contains L4 information.

> **NOTE:** All fragmented packets must be reassembled before the L4 flow can be inferred. This has implications for L4 oriented services like NAT. For NAT to work correctly, the gateway must reassemble every fragmented packet before it can perform its packet transformation. This is a tremendous burden and many NAT implementations (including IPSX) simply do not transform the fragments subsequent to the first.

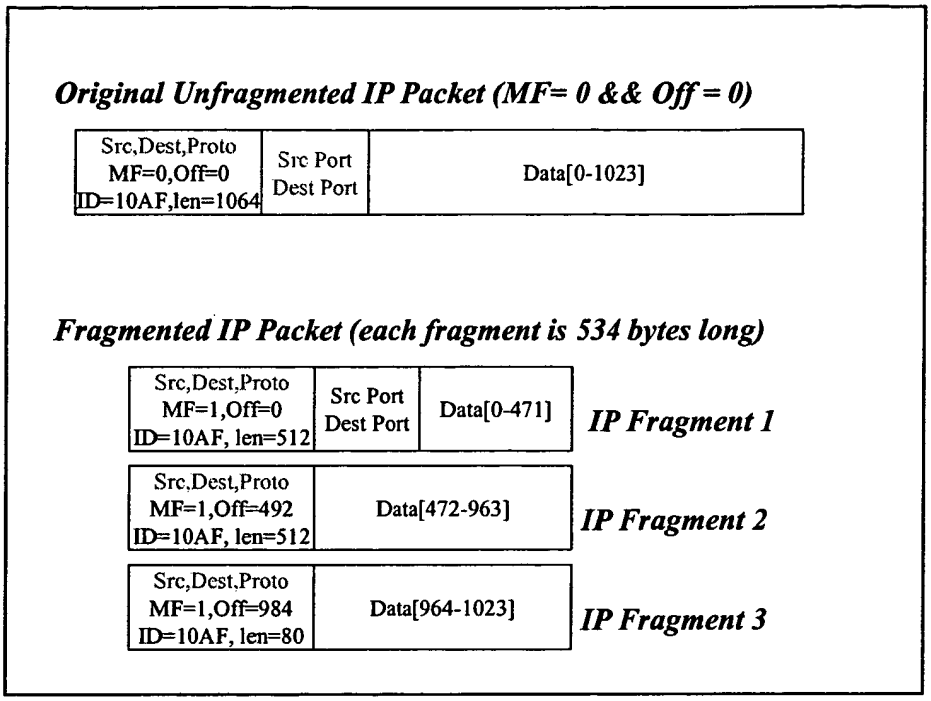

Figure 9: Packet Fragmentation and Header Content

3 Issues and Solutions

3.1 Issues to be addressed

In this section we identify several specific flows that we wish to have optimized. The forward flows are shown as originating at a Subscriber interface and ending at an ISP interface in the upper half of Figure 4, Figure 5, Figure 6 and Figure 7. The reverse flow is shown in the lower half of each figure.

The un-optimized flows are shown in black arrows and represent a fragment of the Configured Topology. The red arrows represent the optimized flows with shortcuts that we would like to generate and represent a fragment of the Flow Topology.

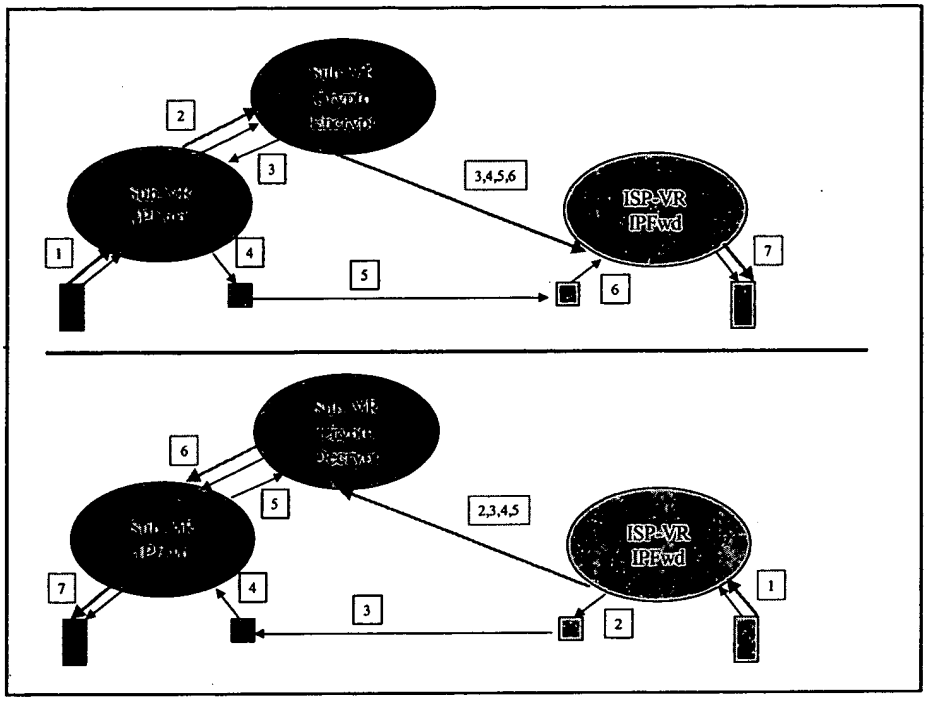

Figure 4: Encryption flows

10

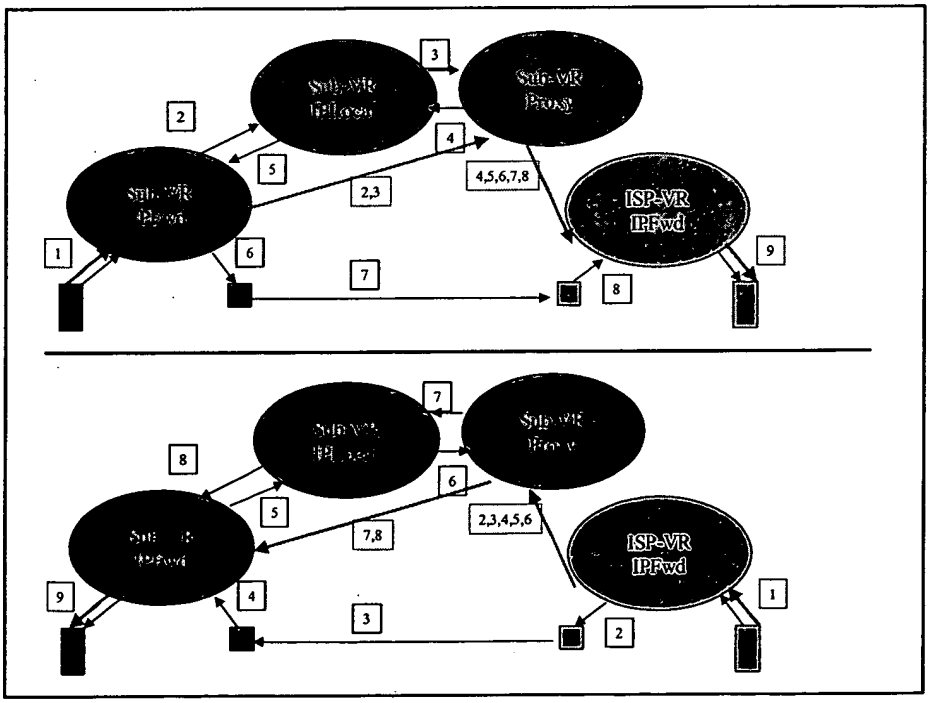

Figure 5: Socket flows

11

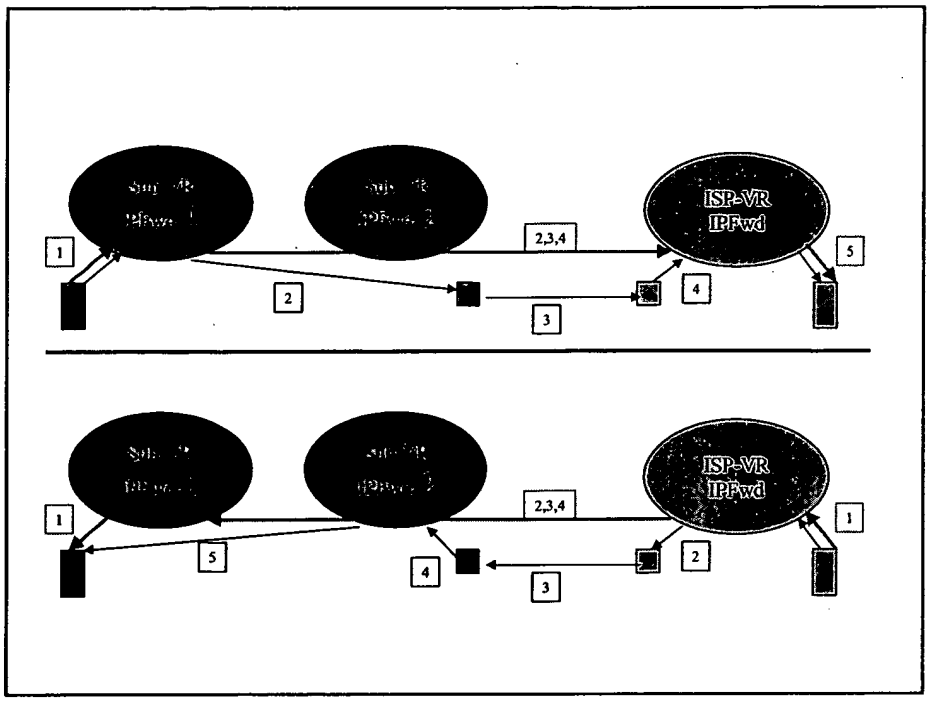

Figure 6: Distributed Forwarding

12

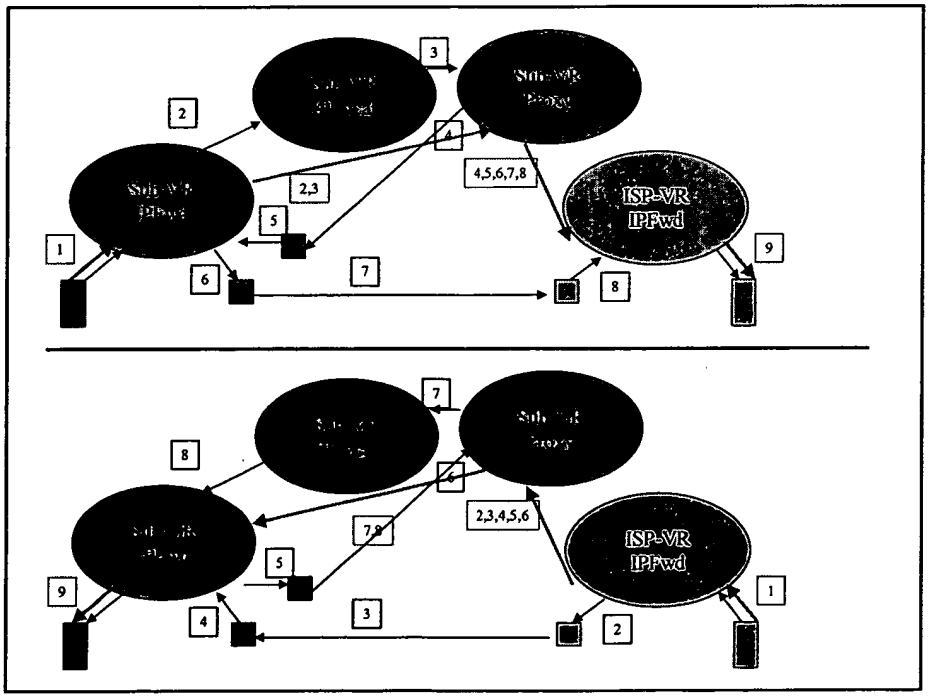

Figure 7: Tunneled flows

13

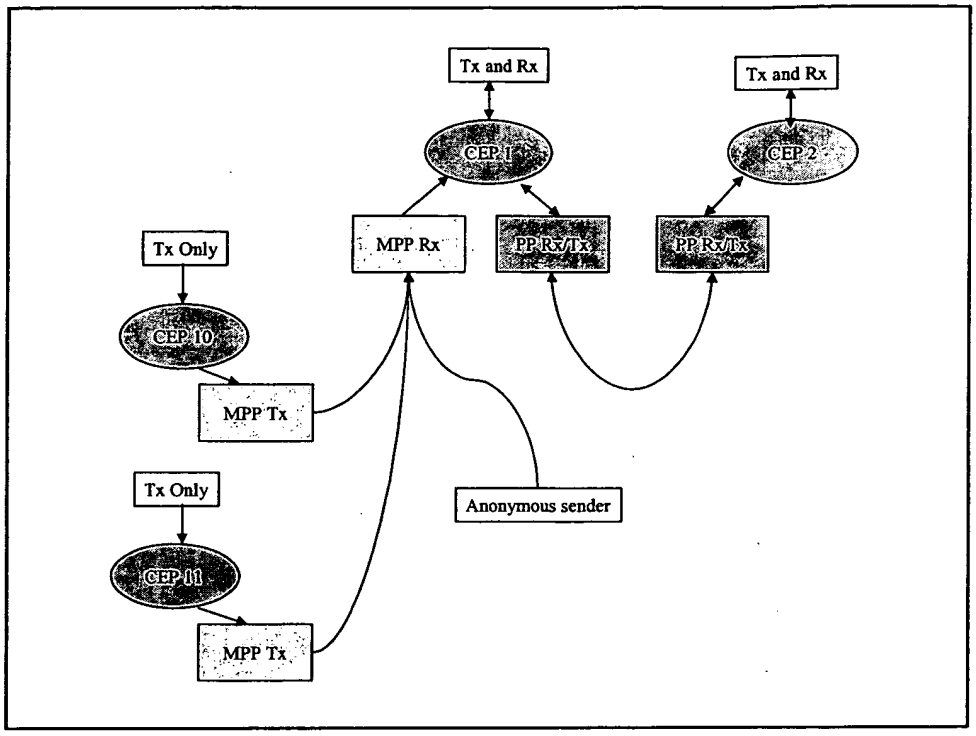

Figure 8: MP-P Operation

3.2.2 Replicated FIB

The IP Forwarding agent uses the Forwarding Information Base (FIB) to make its forwarding decision. When the Virtual Router (VR) terminates links on multiple blades, the VR has to have multiple IP Forwarding agents. This requires the IP Forwarding agents to maintain a Replicated FIB. This feature has been referred to variously as *Multi-FIB*, *Distributed FIB* and *Split FIB* ... none of which reflects the actual operation.

The Replicated FIB need not be constrained to only IP Forwarding agents within a VR. A Replicated FIB may be used by other objects of a VR to make forwarding decisions to support shortcuts.

It should be noted that there is a host other data structures that also need to be replicated. These are

1) Packet Filters (PF), both dynamic and static.
2) Security Policy Database (SPD) and Security Association Database (SAD).
3) Network Address Translation (NAT) tables.

3.2.3 Distributed Protocol Stack (Flow Based)

IPNOS 1.1 supports a functionally distributed protocol stack, which strictly implements the Configured Topology. Figure 2 shows the planned functional decomposition of the

IPSX Virtual Router Objects

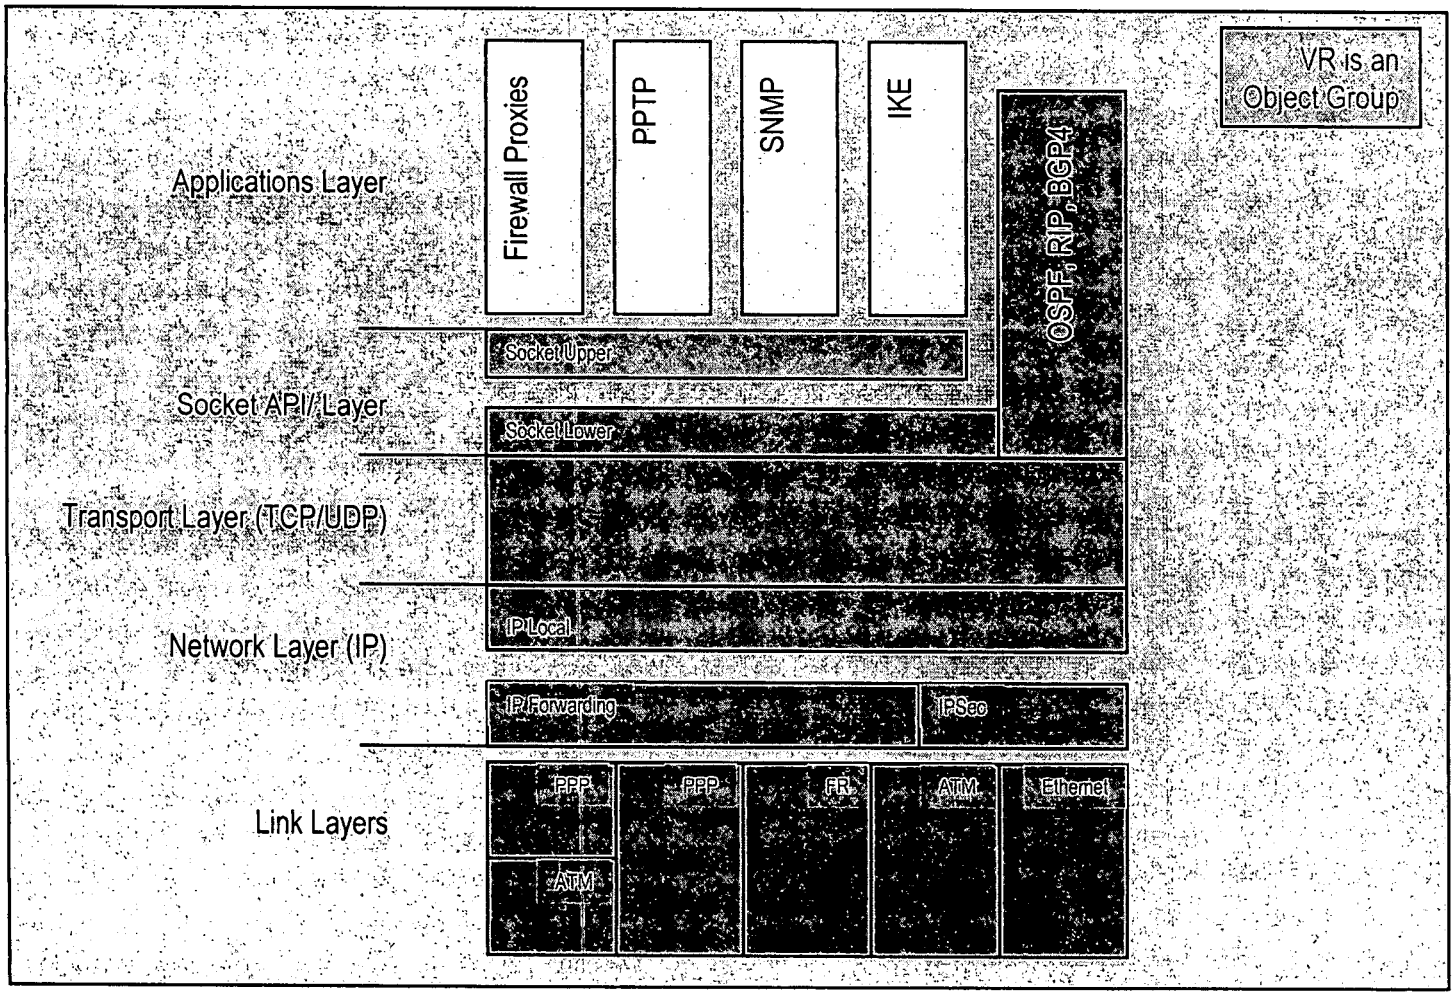

CoSine COMMUNICATIONS

IPSX Virtual Router Object Placement

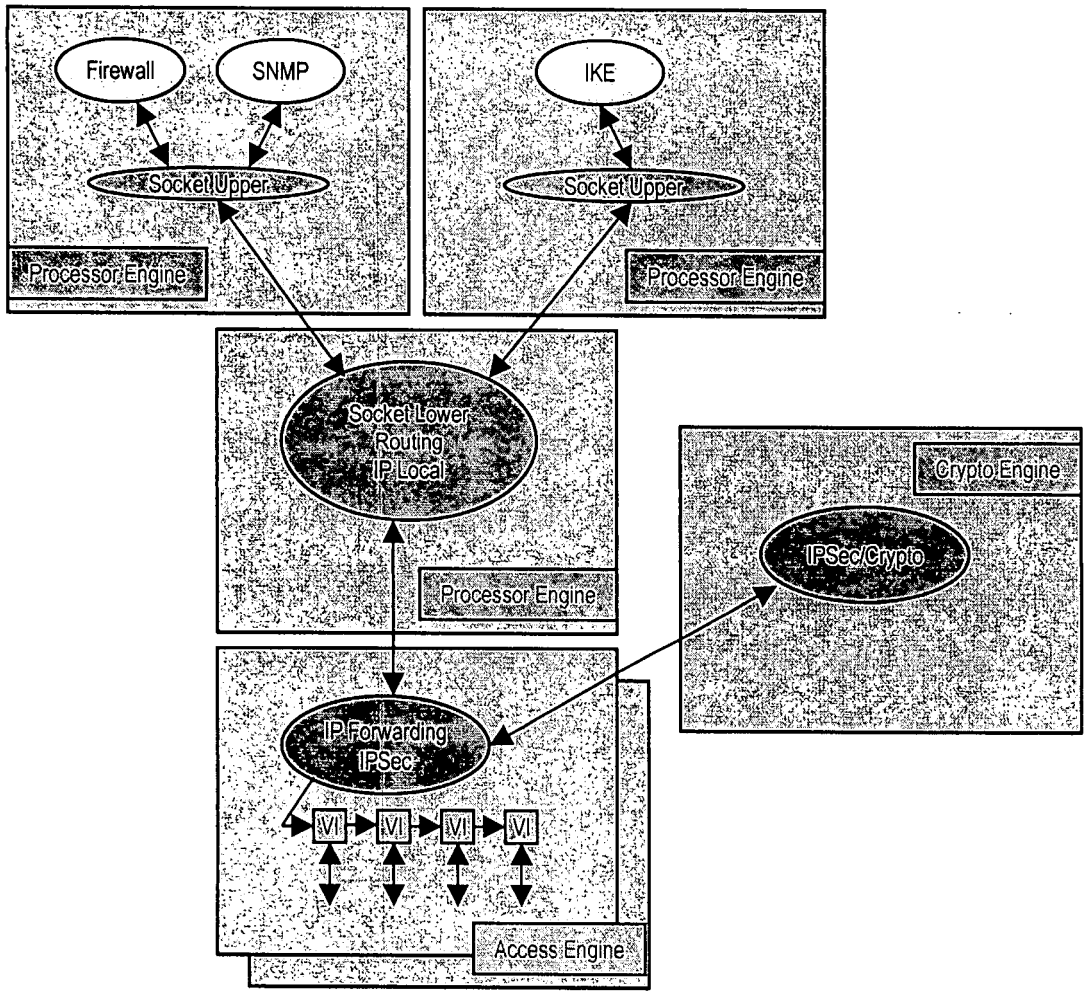

CoSine COMMUNICATIONS

IPSX Virtual Router Instances

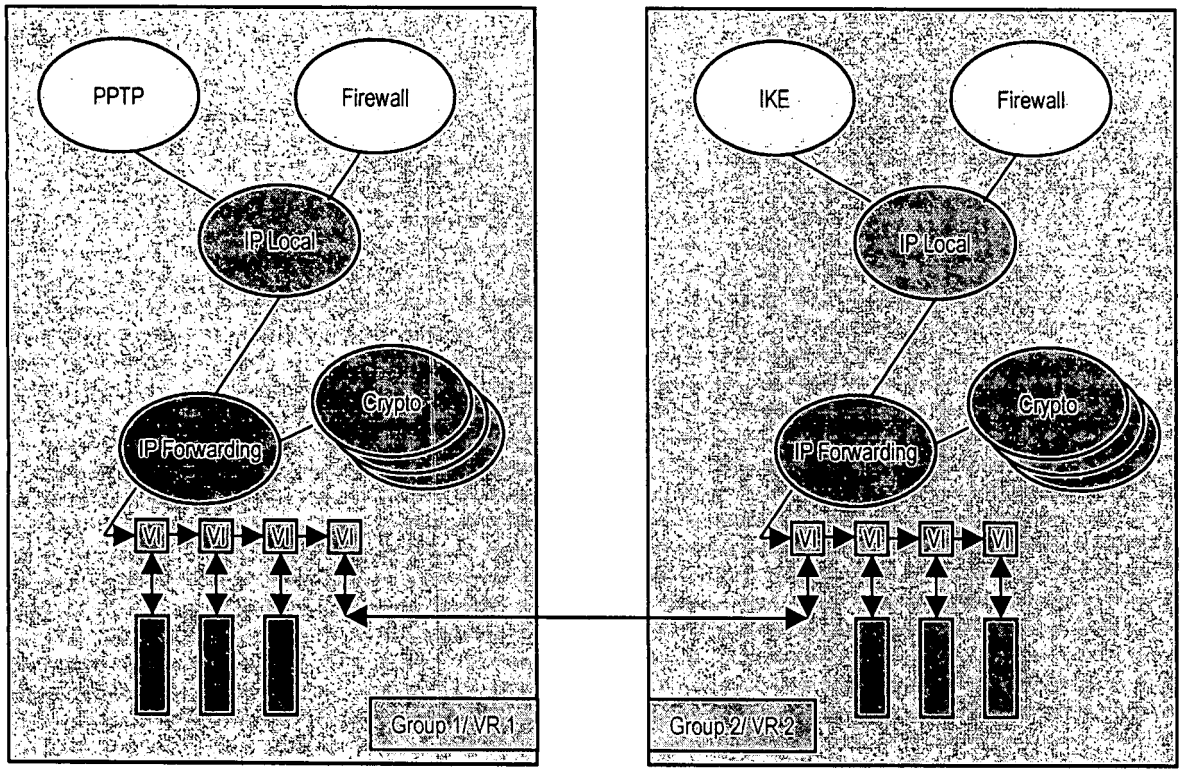

CoSine COMMUNICATIONS

Managed, Service Provider-based VPNs: Intranet Connectivity and Secure Internet Access

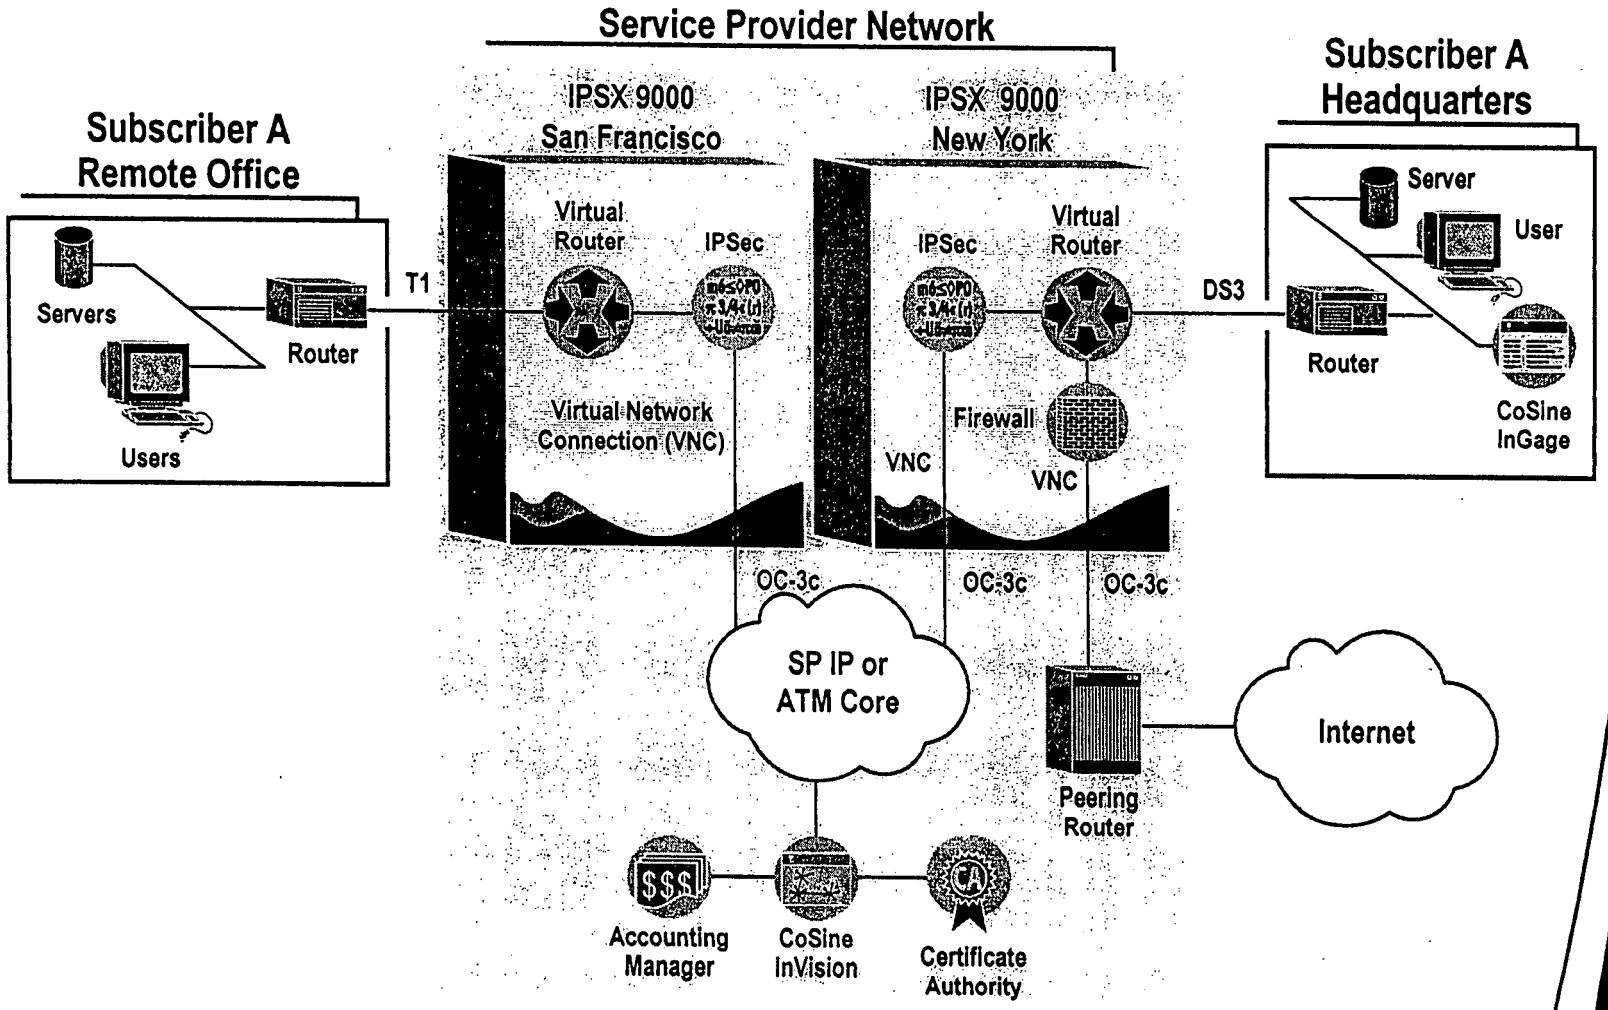

CoSine COMMUNICATIONS

Service Provider Adds Managed Dial-Up Services to VPN

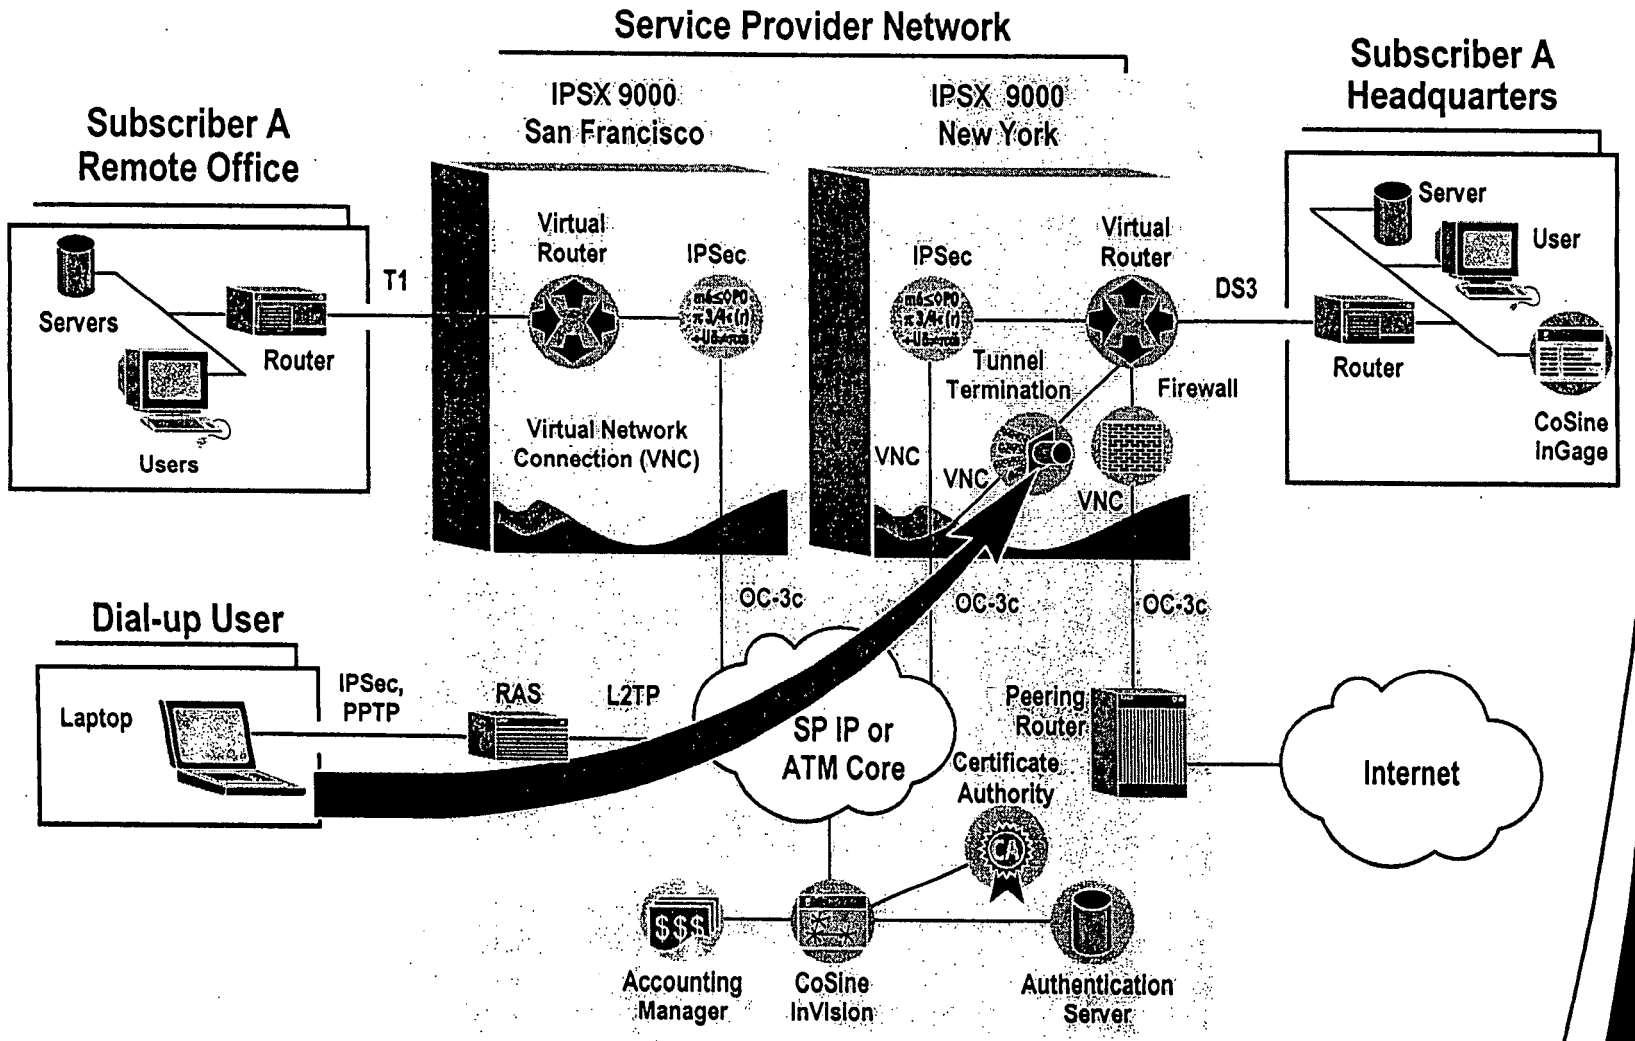

SP Adds Managed Bandwidth Management Services to VPN

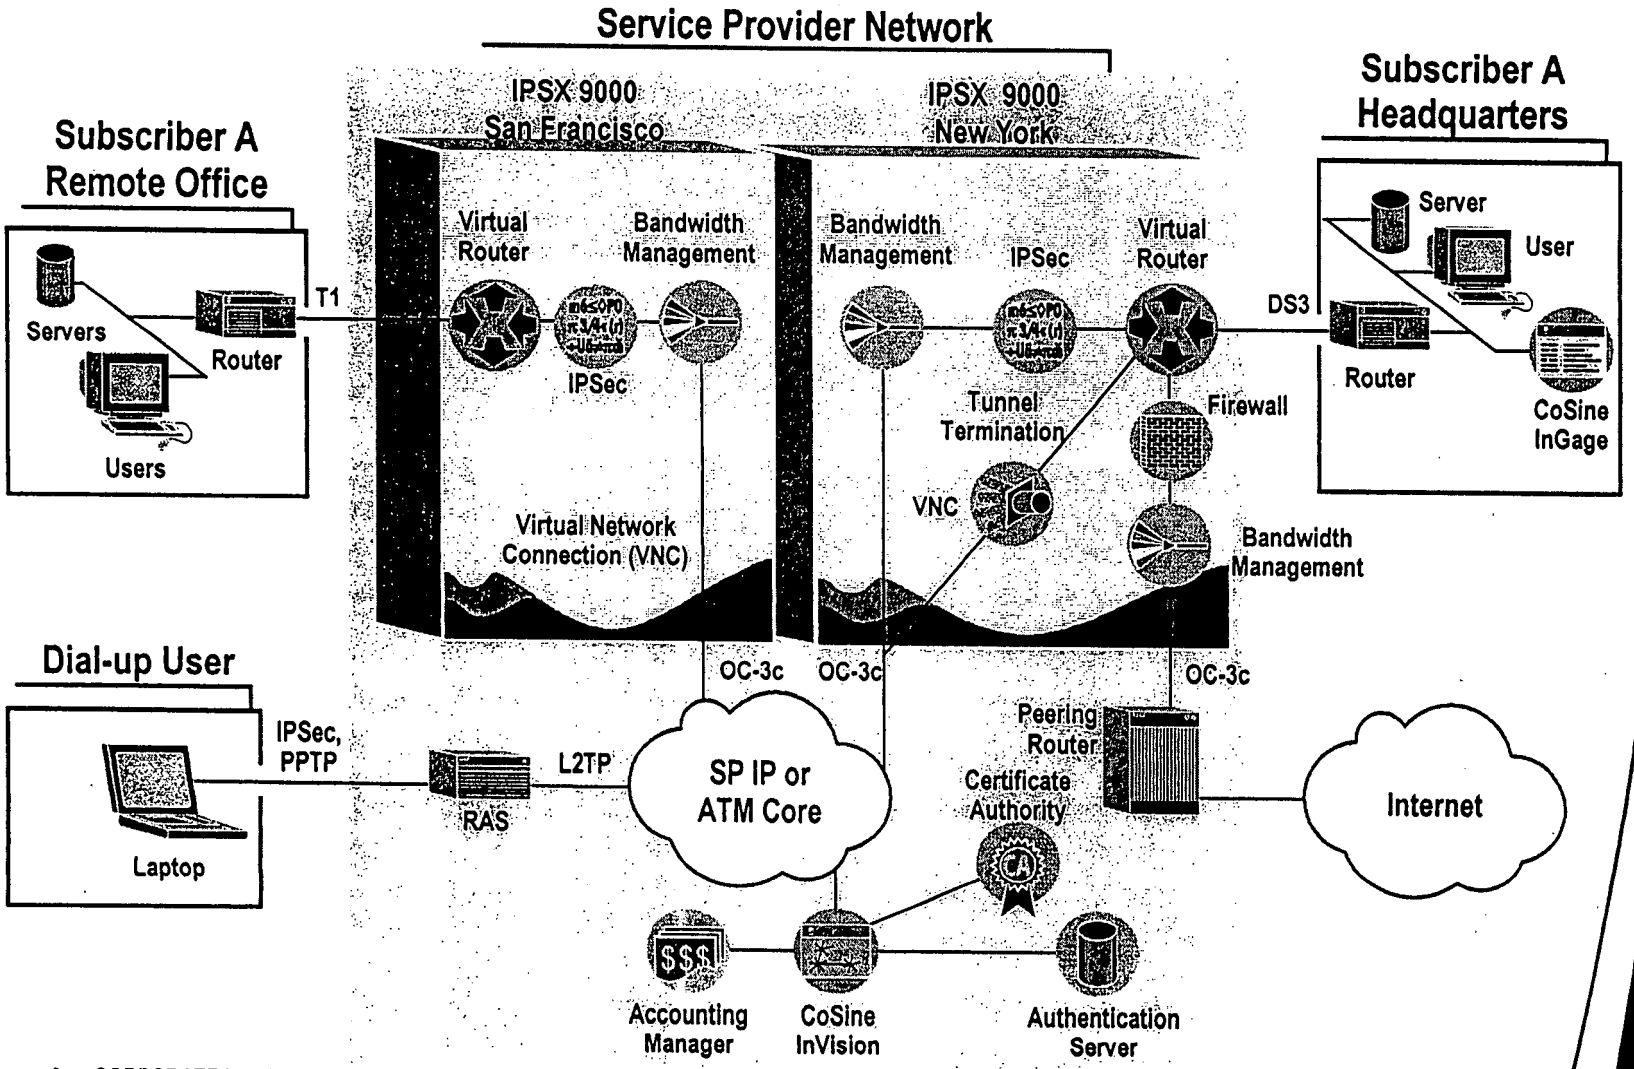

N+I - Atlanta

CoSine COMMUNICATIONS

Dynamic SLAs: CoSine Virtual Network Connection (VNC) Detail

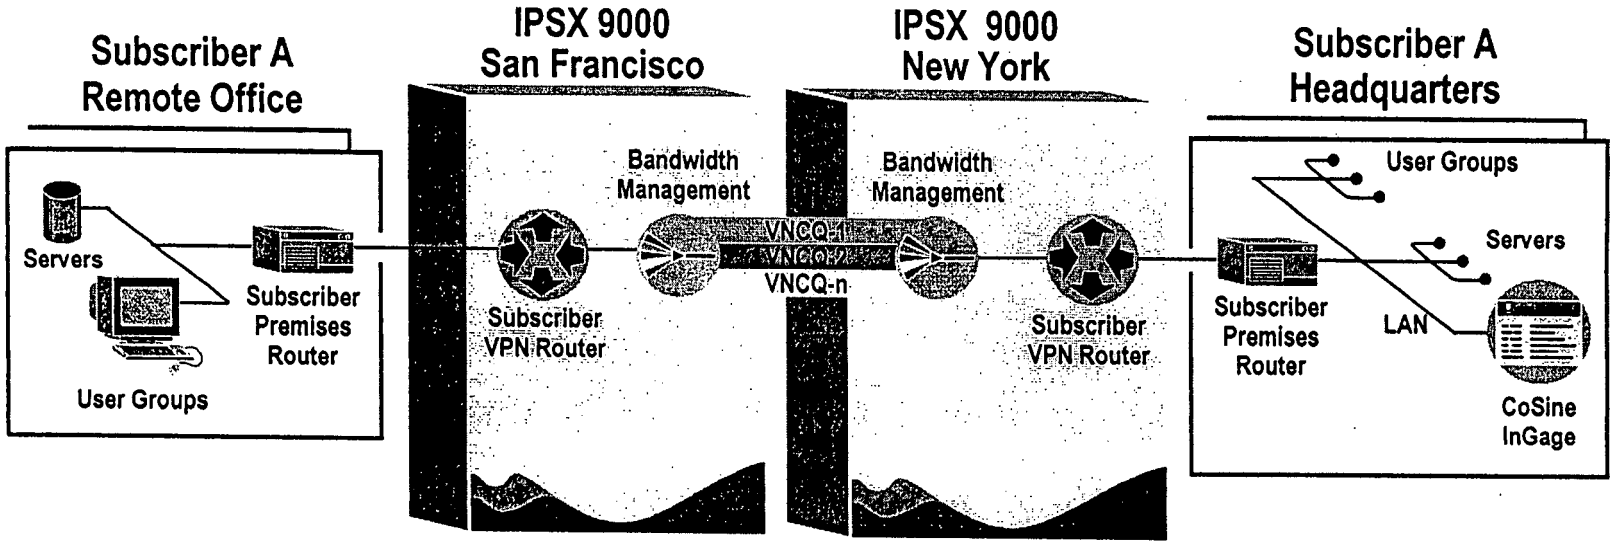

- Bandwidth Management service enables establishment of up to 32 rate shaped queues between virtual routers or on any given interface.
- Each VNCQ may have constituents assigned by user and application, for example:

  VNCQ-1: All engineering FTP traffic – 64 Kbps
  VNCQ-2: All executive officer traffic – 64 Kbps
  VNCQ-n: All http traffic – 12 Kbps

CoSine COMMUNICATIONS

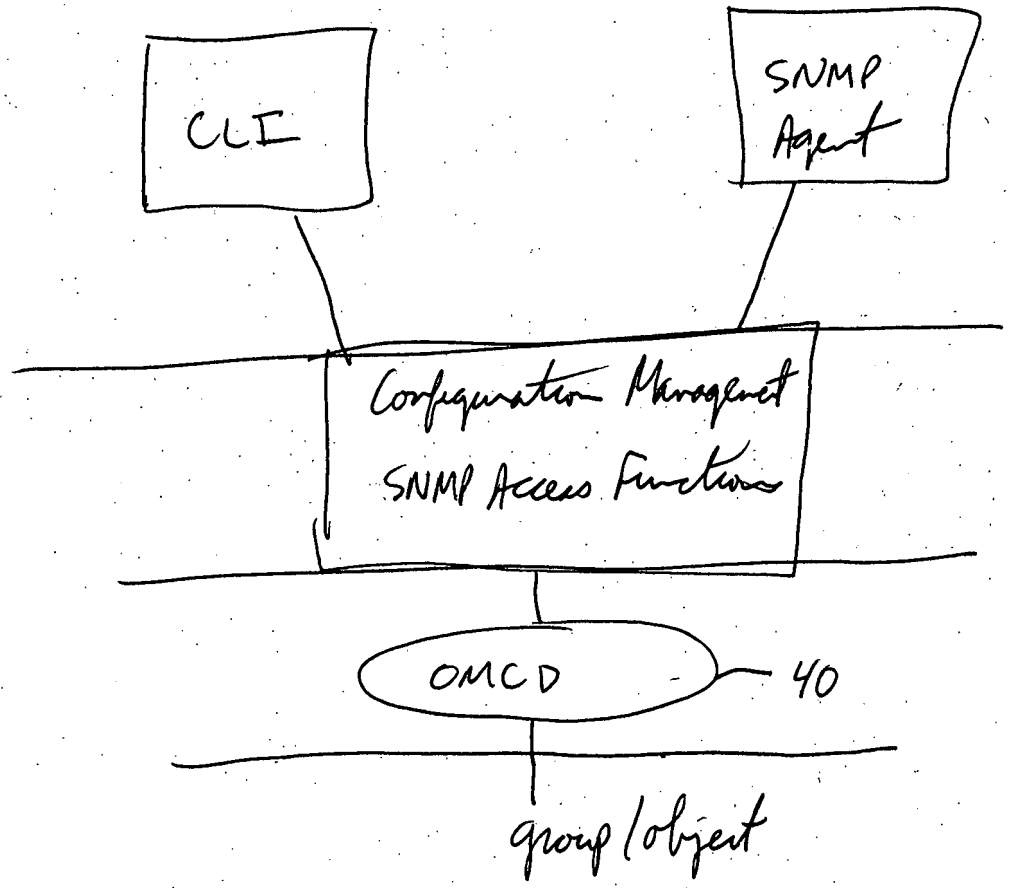
CLI
SNMP Agent
Configuration Management
SNMP Access Function
OMCD
40
group/object

Fig.